Remarks

Reconsideration and allowance of the subject patent application are respectfully requested.

No amendments are made to the claims. A Listing of Claims is provided for the Examiner's convenient reference.

As an initial matter, Applicants respectfully submit that the finality of the 12/15/2006 Office Action is improper. The prior office action does not set forth a proper basis for the rejections of claims 16, 27 and 40 because it fails to address the specific limitations present in these claims. A discussion of the specific limitations of claims 16, 27 and 40 is presented for the first time in the 12/15/2006 Office Action and therefore Applicants believe that the 12/15/2006 Office Action should not have been made final.

Applicants acknowledge with appreciation the indication that claims 13, 15, 24, 26, 36, 38, 41, 42 and 52 contain allowable subject matter.

Applicants respectfully traverse the rejections in the 12/15/2006 Office Action for the reasons set forth in the prior response, which reasons are incorporated herein in their entirety. Applicants provide below comments on the "Response to Argument" section on pages 2-6 of the 12/15/2006 Office Action.

Claims 1-7, 14, 16-18, 25, 27-30, 37, 39, 43-46, 53 and 54 were rejected under 35 U.S.C. Section 103(a) as allegedly being "obvious" over Chin et al. (U.S. Patent No. 6,202,101) in view of Harriman et al. (U.S. Patent No. 6,092,158).

As discussed in the prior response, queues 50a, 50c, 50e, 50d, 50g and 50j of Chin et al. are disclosed therein as being part of interface controller 14, not as being part of the memory controller 44.

The office action responds by alleging:

> ...the memory controller 44 in Chin (*memory controller 44 arbitrates*, Col. 8, lines 56-62) is considered to be equivalent to the arbitration control (825, Figure 8) of the present invention. The interface controller 14 in Chin (*bus interface unit provides an interface between components clocked at similar rates, bus interface unit 14 contains a memory controller,* Col. 7, lines 27-31) is considered to be equivalent to the memory controller (Figure 8) of the present invention. Therefore, the queues in Chin are part of the memory controller. 12/15/2006 Office Action, page 2.

First, such assertions are plainly inconsistent with the Chin et al. disclosure which does not show the queues as being part of memory controller 44. See, e.g., Chin et al., Figure 2.

Second, Applicants respectfully submit that comparisons and conclusory assertions of equivalence between components in Chin et al. and components in Applicants' non-limiting example embodiments (e.g., Figure 8) are not appropriate. Applicants own disclosure cannot be used to re-characterize the disclosure of an applied reference so that this disclosure can thereby be used to reject claims.

Third, Applicants traverse any implication or suggestion in the above excerpt from the office action that "the present invention" is in any way limited to the arrangement shown in Figure 8 of the subject patent application. Figure 8 is a non-limiting example embodiment and the claims are not limited to this specific embodiment.

As discussed in the prior response, the configuration shown in Figure 6 of Chin et al. (including queue 68) is clearly and unambiguously described as being part of processor controller 42, not memory controller 44. See, e.g., Chin et al., col. 11, lines 6-10 ("...Processor controller 42 includes an in-order queue 42, a peripheral request queue 66, and memory request queue 68.") (emphasis added). Consequently, queue 68 cannot constitute the claimed multiple resource buffer which is specified to be part of the memory controller.

The office action responds by alleging:

> Chin teaches that the processor controller 42 is part of the interface controller 14 (Col. 8, lines 23-26), which is equivalent to the memory controller, as discussed above. Therefore, queue 68 is part of the memory controller. 12/15/2006 Office Action, page 2.

Figure 2 of Chin et al. clearly shows that the processor controller 42 and memory controller 44 are distinct components of interface controller 14 and Figure 6 of Chin et al. clearly shows that queue 68 is part of processor controller 42, not memory controller 44. An assertion that the entirety of the interface controller 14 can be viewed as a memory controller is simply a contrivance to remedy Chin et al.'s manifest failure to disclose that queue 68 is part of memory controller 44.

As discussed in the prior response, Chin et al. describes queue 68 as storing requests received from a processor bus and thus queue 68 cannot constitute the claimed multiple resource buffer which is specified as storing requests for memory access from a plurality of buffer memories for multiple resources.

The office action responds by alleging:

> Chin teaches that memory request queue 68 is coupled to the plurality of buffer memories (queues) for storing requests for main memory access from each of the plurality of resources (processors) ... as recited in the claims, and therefore suggests a control circuit involved in transferring information. 12/15/2006 Office Action, page 3.

The assertion that queue 68 stores requests for main memory access from buffer memories from multiple resources is simply incorrect. Even assuming that queue 68 is viewed as being connected to other queues, "[q]ueues 64, 66, and 68 receive various items of information associated with a request transaction dispatched on the processor bus." Chin et al., col. 11, lines 11-13 (emphasis added). Chin et al.'s description of Figure 5 confirms that queue 68 receives requests from the processor bus. See, e.g., col. 12, lines 44-47 ("Referring to FIG. 5, details of information which can be stored in queues 64, 66, and 68 are shown. Further shown is an example of a sequence of requests forwarded from the processor bus.") While the requests in queue 68 from the processor bus might be sent to different destinations (e.g., memory or peripheral), the requests are from a single resource and there is no disclosure or suggestion in Chin et al. that queue 68 is a multiple resource buffer for storing requests for main memory access from each of a plurality of resources.

As discussed in the prior response, queue 50c is unrelated to the concept of controlling a transfer of information from buffer memories to a multiple resource buffer as claimed. This is confirmed by reference to Table I of Chin et al. which shows that M2P queue 50c contains data whose source is the memory 18 and whose destination is processor 12. The 12/15/2006 office action offers no response to this point. The remaining disclosure of Chin et al. is consistent with this Table I entry for M2P and confirms that queue 50c is not involved with a transfer of information from buffer memories to a multiple resource buffer.

The office action discussion regarding queue 50c references Chin et al, col. 12, lines 65-67: "If the request [forwarded from the processor bus] is a read request to memory, the memory request will access a memory location and data at that location is temporarily stored within an M2P queue 50c." Data dequeued from queue 50c is presented to the processor bus. See, e.g., Chin et al., col. 13, lines 6-7. These portions of Chin et al. teach that data read from the memory in response to a read request from the processor is queued in queue 50c and is thereafter "presented" to the processor. There is no concept of a multiple resource buffer here.

Queue 50c is also mentioned at col. 13, lines 27 et seq. in the discussion of Figure 6 (which illustrates an alternative configuration to that of Figure 5). This portion of Chin et al. describes: "More specifically, a block diagram illustrates the order of requests associated with a memory read or write are maintained in that same order when the corresponding data is placed in memory data queue 50c/50a." Thus, Chin et al. describes that the order of read requests from the processor is maintained when the corresponding data is retrieved from the memory and placed in queue 50c. Similarly, the order of write requests from the processor is maintained when the corresponding data is placed in queue 50a. Here again, there is no concept of a multiple resource buffer.

The discussion bridging pages 3 and 4 of the 12/152006 Office Action is fundamentally flawed because queue 50c is apparently viewed in the office action as a "request queue." As is evident from the discussion above, queue 50c stores data read from the memory after a memory request is processed. Queue 50c is not a request queue.

The office action also references the col. 13, line 51 – col. 14, line 5 disclosure of Chin et al. in connection with the claimed control circuit. As previously discussed, this disclosure is unrelated to the claimed control circuit. Instead, this disclosure describes queues within processor controller 42 – queues that involve requests from the processor bus, not from a plurality of buffer memories each of which is operatively coupled to one of a plurality of resources requesting memory access.

The office action cites Harriman et al. for its reference to the grouping of reads and writes in order to "reduce turnaround." Harriman et al., col. 1, lines 43-47.

Applicants respectfully submit that Harriman et al. does not remedy the deficiencies of Chin et al. discussed above with respect to the claimed buffer memories, multiple resource buffer memory and/or control circuit. As such, even assuming for the sake of argument that Harriman et al.'s technique of reducing turnaround is viewed as reducing the switching frequency of main memory read and write operations and this technique were forcibly combined with Chin et al., the subject matter of claim 1 could still not possibly result.

For at least these reasons, Applicants respectfully submit that the subject matter of claim 1 and its dependent claims 2-7 and 14 would not have been made obvious by the proposed combination of Chin et al. and Harriman et al.

These dependent claims contain features that provide additional bases for patentability.

By way of example, claim 14 requires that a resource that is writing to main memory generates a flush signal for initiating the flushing of that resource's write request queue. The office action contends that the "de-queuing" described at col. 13, lines 1-11 of Chin et al. discloses this feature. However, this de-queuing is for queue 50c, which contains data whose source is the memory and whose destination is the processor. This queue is not a write request queue and there is no disclosure or suggestion in Chin et al. of initiating the flushing of a resource's write request queue when that resource is writing to main memory. For this additional and independent reason, claim 14 is believed to distinguish over the proposed combination of Chin et al. and Harriman et al.

The discussion of claim 14 in the 12/152006 Office Action contends that queue 50c is a "write queue." This is not the case. The text of the office action references col. 12, lines 65-67 for its disclosure that the memory request will access a memory location and data at that location is temporarily stored within an M2P queue. The office action however omits the first few words of col. 12, lines 65-67 which state "[i]f the request is a read request to memory" (emphasis added). Queue 50c stores data retrieved from the memory and is simply not a write queue, contrary to the assertion in the office action.

Claim 16 is for a memory controller comprising a main processor related interface including read and write request queues; a first resource related interface including read and write request queues; a second resource related interface including read and write queues; and a multiple resource write request queue. Among other things, this rejection is premised in part on the identification of queue 50c as a "read request queue". As discussed above, queue 50c is not a read request queue and thus the rejection fails to identify in Chin et al. a read request queue for a main processor related interface.

Moreover, as further discussed above, queue 68 is not a multiple resource write request queue since queue 68 is only described in Chin et al. as storing requests from the processor.

Harriman et al. does not remedy the deficiencies of Chin et al. with respect to these features and, consequently, the proposed combination of these documents is deficient in this regard with respect to claim 16 and its dependent claims 17, 18 and 25.

These dependent claims contain features that provide additional bases for patentability.

By way of example without limitation, claim 25 requires that a resource that is writing to main memory generates a flush signal for initiating the flushing of that resource's write request queue. The office action contends that the "de-queuing" described at col. 13, lines 1-11 of Chin

et al. discloses this feature. However, as noted above in the discussion of claim 14, this de-queuing operation is for queue 50c, which contains data whose source is the memory and whose destination is the processor. This queue is not a write request queue and there is no disclosure or suggestion in Chin et al. of initiating the flushing of a resource's write request queue when that resource is writing to main memory. For this additional and independent reason, claim 25 is believed to distinguish over the proposed combination of Chin et al. and Harriman et al.

Claim 27 is for a method of controlling access to a main memory which includes "delaying forwarding requests for main memory access to a memory access control circuit to reduce the frequency of switching between memory read states and memory write states" and "granting requests for main memory access by said memory access control circuit." The office action acknowledges that Chin et al. does not teach the delaying step, but references Harriman as disclosing the "grouping" of read and writes to reduce turn around time. According to the office action, this "grouping" is "similar" to the claimed "delaying". The office action provides no explanation as to how these concepts are purportedly "similar" and Applicants respectfully submit that a conclusory allegation of "similarity" does not constitute evidence to support the legal conclusion of obviousness.

Applicants also note that queues 50d and 50h are identified in the discussion of request queues in the rejection of claim 27. However, these queues are in fact a memory-to-PCI queue and a memory-to-AGP queue, respectively, and store data read from the memory, not requests for such data.

Consequently, the proposed combination of Chin et al. and Harriman et al. would not have resulted in the subject matter of claim 27 or its dependent claims 28-30, 37, 39, 40, 43-46 and 53.

These dependent claims contain features that provide additional bases for patentability.

By way of example without limitation, claim 37 requires that a resource that is writing to main memory generates a flush signal for initiating the flushing of that resource's write request queue. The office action contends that the "de-queuing" described at col. 13, lines 1-11 of Chin et al. discloses this feature. However, as noted above in the discussion of claim 14, this de-queuing is with respect to queue 50c, which contains data whose source is the memory and whose destination is the processor. This queue is not a write request queue and there is no disclosure or suggestion in Chin et al. of initiating the flushing of a resource's write request queue when that resource is writing to main memory. For this additional and independent

reason, claim 37 is believed to distinguish over the proposed combination of Chin et al. and Harriman et al.

Claim 40 is for a method of controlling access to a main memory which requires generating a write queue flush signal by a first resource to initiate copying information in the first resource write request queue to main memory and flushing the first resource write request queue. Among other things, the col. 13, lines 1-11 disclosure of Chin et al. referenced in the office action in connection with these features does not describe or suggest generating a flush signal or flushing a write request queue. The first lines of this disclosure describe operations involving data read from the memory and stored in queue 50c. There is nothing here about write queues or the flushing of such queues. The last few lines of this disclosure state: "If the memory request is a write request, then the address will be held in the P2M queue (queue 50a) until that request's entry number matches the current in-order queue entry number." This description relates to write requests, but simply describes how long an address is held in a queue. There is nothing about the flushing of a write request queue. Consequently, claim 40 and its dependent claims 43-51 are not made obvious by the proposed combination of Chin et al. and Harriman et al.

Like claim 1, claim 54 similarly calls for buffer memories, a multiple resource buffer and control circuit. Consequently, Applicants respectfully submit that the subject matter of claim 54 would not have been made obvious by the proposed combination of Chin et al. and Harriman et al.

Claims 8-12, 19-23, 31-35 and 47-51 were rejected under 35 U.S.C. Section 103(a) as allegedly being "obvious" over the proposed Chin et al-Harriman et al. combination, in view of Jeddeloh et al. (U.S. Patent No. 6,330,647). Applicants respectfully traverse this rejection.

Jeddeloh et al. was applied in connection with its alleged disclosure of an arbiter and control registers. However, even assuming that such features were somehow shown to be properly combinable with the result of the Chin et al.-Harriman et al. combination, Jeddeloh et al. does not remedy the deficiencies of Chin et al. and Harriman et al. with respect to claims 1, 16, 27 and 40, from which claims 8-12, 19-23, 31-35 and 47-51 depend.

The pending claims are believed to be in condition for allowance and favorable office action is respectfully requested.

Respectfully submitted,

NIXON & VANDERHYE P.C.

By: 

Michael J. Shea
Reg. No. 34,725

MJS:mjs
901 North Glebe Road, 11th Floor
Arlington, VA 22203-1808
Telephone: (703) 816-4000
Facsimile: (703) 816-4100